arXiv:1609.09408v1 [stat.ML] 29 Sep 2016

# Cooperative Training of Descriptor and Generator Networks

**Jianwen Xie** JIANWEN@UCLA.EDU
**Yang Lu** YANGLV@UCLA.EDU
**Song-Chun Zhu** SCZHU@STAT.UCLA.EDU
**Ying Nian Wu** YWU@STAT.UCLA.EDU
Department of Statistics, University of California, Los Angeles, CA, USA

## Abstract

This paper studies the cooperative training of two probabilistic models of signals such as images. Both models are parametrized by convolutional neural networks (ConvNets). The first network is a descriptor network, which is an exponential family model or an energy-based model, whose feature statistics or energy function are defined by a bottom-up ConvNet, which maps the observed signal to the feature statistics. The second network is a generator network, which is a non-linear version of factor analysis. It is defined by a top-down ConvNet, which maps the latent factors to the observed signal. The maximum likelihood training algorithms of both the descriptor net and the generator net are in the form of alternating back-propagation, and both algorithms involve Langevin sampling. We observe that the two training algorithms can cooperate with each other by jumpstarting each other's Langevin sampling, and they can be naturally and seamlessly interwoven into a CoopNets algorithm that can train both nets simultaneously.

## 1. Introduction

### 1.1. Two ConvNets of opposite directions

We begin with a tale that the reader of this paper can readily relate to. A student writes up an initial draft of a paper. His advisor then revises it. After that they submit the revised paper for review. The student then learns from his advisor's revision, while the advisor learns from the outside review. In this tale, the advisor guides the student, but the student does most of the work.

Such is a tale we are going to tell in this paper. It is a tale of two nets. It is about two probabilistic models of signals such as images. Both models are parametrized by convolutional neural networks (ConvNets or CNNs) (LeCun et al., 1998; Krizhevsky et al., 2012). As mighty as ConvNets have become, they are still one way streets. So the two nets in our tale take two opposite directions. One is bottom-up, and the other is top-down, as illustrate by the following diagram:

$$\begin{array}{cc} \text{Bottom-up ConvNet} & \text{Top-down ConvNet} \\ \textbf{features} & \textbf{latent variables} \\ \Uparrow & \Downarrow \\ \text{signal} & \text{signal} \\ \text{(a) Descriptor Net} & \text{(b) Generator Net} \end{array} \tag{1}$$

The simultaneous training of such two nets was first studied by the recent work of (Kim & Bengio, 2016). These two nets belong to two major classes of probabilistic models. (a) The exponential family models or the energy-based models (LeCun et al., 2006; Hinton, 2002) or the Markov random field models (Zhu et al., 1997; Roth & Black, 2005), where the probability distribution is defined by the feature statistics or the energy function computed from the signal by a bottom-up process. (b) The latent variable models or the directed graphical models, where the signal is assumed to be a transformation of the latent factors that follow a known prior distribution. The latent factors generate the signal by a top-down process. A classical example is factor analysis (Rubin & Thayer, 1982).

The two classes of models have been contrasted by (Zhu, 2003; Guo et al., 2003; Teh et al., 2003; Wu et al., 2008; Ngiam et al., 2011). (Zhu, 2003; Guo et al., 2003) called the two classes of models the descriptive models and the generative models respectively. Both classes of models can benefit from the high capacity of the multi-layer ConvNets. (a) In the exponential family models or the energy-based models, the feature statistics or the energy function can be defined by a bottom-up ConvNet that maps the signal to the features (Ngiam et al., 2011; Dai et al., 2015; Lu et al., 2016; Xie et al., 2016), and the energy function is usually the sum or a linear combination of the features at the top layer. We call the resulting model a descriptive network or

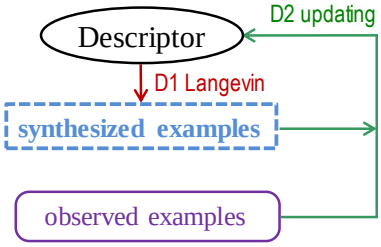


*Figure 1.* The flow chart of Algorithm D for training the descriptor net. The updating in Step D2 is based on the difference between the observed examples and the synthesized examples. The Langevin sampling of the synthesized examples from the current model in Step D1 can be time consuming.

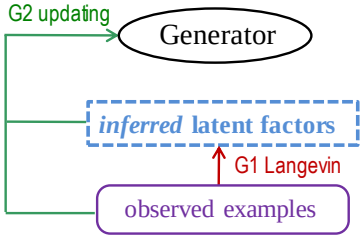


*Figure 2.* The flow chart of Algorithm G for training the generator net. The updating in Step G2 is based on the observed examples and their inferred latent factors. The Langevin sampling of the latent factors from the current posterior distribution in Step G1 can be time consuming.

a descriptor net following (Zhu, 2003), because it is built on descriptive feature statistics. (b) In the latent variable models or the directed graphical models, the transformation from the latent factors to the signal can be defined by a top-down ConvNet (Zeiler & Fergus, 2014; Dosovitskiy et al., 2015), which maps the latent factors to the signal. We call the resulting model a generative network or generator net following (Goodfellow et al., 2014), who proposed such a model in their work on the generative adversarial networks (GAN).

In this article, we assume that the two nets have separate structures and parameters. What we shall do is to fuse the maximum likelihood training algorithms of these two nets. Although each training algorithm can stand on its own, their fusion is both beneficial and natural, in particular, the fusion fuels the Markov chain Monte Carlo (MCMC) for sampling the descriptor net, and turns the training of the generator net from unsupervised learning into supervised learning. This fusion leads to a cooperative training algorithm that interweaves the steps of the two algorithms seamlessly. We call the resulting algorithm the CoopNets algorithm, and we show that it can train both nets simultaneously.

### 1.2. Two algorithms of alternating back-propagation

Both the descriptor net and the generator net can be learned from the training examples by maximum likelihood.

The training algorithm for the descriptor net alternates between the following two steps (Xie et al., 2016). We call it Algorithm D. See Figure 1 for an illustration.

**Step D1 for Langevin revision**: Sampling synthesized examples from the current model by Langevin dynamics (Neal, 2011; Girolami & Calderhead, 2011). We call this step Langevin revision because it keeps revising the current synthesized examples.

**Step D2 for density shifting**: Updating the parameters of the descriptor net based on the difference between the observed examples and the synthesized examples obtained in D1. This step is to shift the high probability regions of the descriptor from the synthesized examples to the observed examples.

Both steps can be powered by back-propagation. The algorithm is thus an alternating back-propagation algorithm.

In Step D1, the descriptor net is dreaming by synthesizing images from the current model. In Step D2, the descriptor updates its parameters to make the dream more realistic.

The training algorithm for the generator net alternates between the following two steps (Han et al., 2016). We call it Algorithm G. See Figure 2 for an illustration.

**Step G1 for Langevin inference**: For each observed example, sample the latent factors from the current posterior distribution by Langevin dynamics. We call this step Langevin inference because it infers the latent factors for each observed example, in order for the inferred latent variables to explain or reconstruct the observed examples.

**Step G2 for reconstruction**: Updating the parameters of the generator net based on the observed examples and their inferred latent factors obtained in G1, so that the inferred latent variables can better reconstruct the observed examples.

Both steps can be powered by back-propagation. The algorithm is thus an alternating back-propagation algorithm.

The training algorithm for generator net is similar to the EM algorithm (Dempster et al., 1977), where step G1 can be mapped to the E-step, and step G2 can be mapped to the M-step. It is an unsupervised learning algorithm because the latent factors are unobserved.

In Step G1, the generator net is thinking about each observed example by inferring the latent factors that can reconstruct it. The thinking involves explaining-away rea-

soning: the latent factors compete with each other in the Langevin inference process to explain the example. In Step G2, the generator net updates its parameters to make the thinking more accurate.

Compared to Step D2, Step G2 is also a form of density shifting from the current reconstructed examples towards the observed examples, but it requires inferring the latent factors for each observed example.

### 1.3. CoopNets algorithm: reconstruct the revisions

The two algorithms can operate separately on their own (Xie et al., 2016; Han et al., 2016). But just like the case with the student and the advisor, they benefit from cooperating with each other, where the generator net plays the role of the student, and the descriptor net plays the role of the advisor.

The needs for cooperation stem from the fact that the Langevin sampling in Step D1 and Step G1 can be time-consuming. If the two algorithms cooperate, they can jumpstart each other's Langevin sampling in D1 and G1. The resulting algorithm, which we call the CoopNets algorithm, naturally and seamlessly interweaves the steps in the two algorithms with minimal modifications.

In more plain language, while the descriptor needs to dream hard in Step D1 for synthesis, the generator needs to think hard in Step G1 for explaining-away reasoning. On the other hand, the generator is actually a much better dreamer because it can generate images by direct ancestral sampling, while the descriptor does not need to think in order to learn.

Specifically, we have the following steps in the CoopNets algorithm:

**Step G0 for generation**: Generate the initial synthesized examples using the generator net. These initial examples can be obtained by direct ancestral sampling.

**Step D1 for Langevin revision**: Starting from the initial synthesized examples produced in Step G0, run Langevin revision dynamics for a number of steps to obtain the revised synthesized examples.

**Step D2 for density shifting**: The same as before, except that we use the revised synthesized examples produced by Step D1 to shift the density of the descriptor towards the observed examples.

**Step G1 for Langevin inference**: The generator net can learn from the revised synthesized examples produced by Step D1. For each revised synthesized example, we know the values of the latent factors that generate the corresponding initial synthesized example in Step G0, therefore we may simply infer the latent factors to be their known values given in Step G0, or initialize the Langevin inference dynamics in Step G1 from the known values.

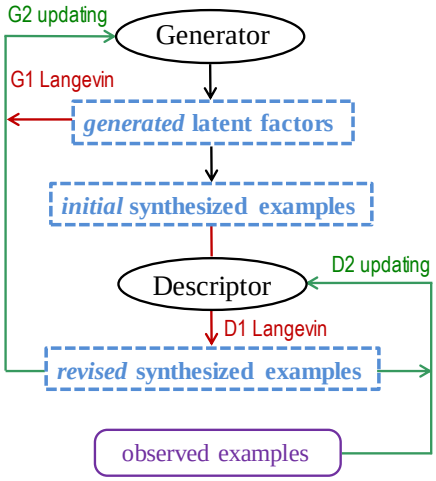


*Figure 3.* The flow chart of the CoopNets algorithm. The part of the flow chart for training the descriptor is similar to Algorithm D in Figure 1, except that the D1 Langevin sampling is initialized from the initial synthesized examples supplied by the generator. The part of the flow chart for training the generator can also be mapped to Algorithm G in Figure 2, except that the revised synthesized examples play the role of the observed examples, and the known generated latent factors can be used as inferred latent factors (or be used to initialize the G1 Langevin sampling of the latent factors).

**Step G2 for reconstruction**: The same as before, except that we use the revised synthesized examples and the inferred latent factors obtained in Step G1 to update the generator. The generator in Step G0 generates and thus reconstructs the initial synthesized examples. Step G2 updates the generator to reconstruct the revised synthesized examples. The revision of the generator accounts for the revisions made by the Langevin revision dynamics in Step D1.

Figure 3 shows the flow chart of the CoopNets algorithm. The generator is like the student. It generates the initial draft of the synthesized examples. The descriptor is like the advisor. It revises the initial draft by running a number of Langevin revisions. The descriptor learns from the outside review, which is in the form of the difference between the observed examples and the revised synthesized examples. The generator learns from how the descriptor revises the initial draft by reconstructing the revised draft. Such is the tale of two nets.

The reason we let the generator net learn from the revised synthesized examples instead of the observed examples is that the generator does not know the latent factors that generate the observed examples, and it has to think hard to infer them by explaining-away reasoning. However, the

generator knows the latent factors that generate each initial synthesized example, and thus it essentially knows the latent factors when learning from the revised synthesized examples. By reconstructing the revised synthesized examples, the generator traces and accumulates the Langevin revisions made by the descriptor. This cooperation is thus beneficial to the generator by relieving it the burden of inferring the latent variables.

While the generator may find it hard to learn from the observed examples directly, the descriptor has no problem learning from the observed examples because it only needs to compute the bottom-up features deterministically. However, it needs synthesized examples to find its way to shift its density, and they do not come by easily. The generator can provide unlimited number of examples, and in each learning iteration, the generator supplies a completely new batch of independent examples on demand. The descriptor only needs to revise the new batch of examples instead of generating a new batch by itself from scratch. The generator has memorized the cumulative effect of all the past Langevin revisions, so that it can produce new samples in one shot. This cooperation is thus beneficial to the descriptor by relieving it the burden of synthesizing examples from scratch.

In terms of density shifting, the descriptor shifts its density from the revised synthesized examples towards the observed examples in Step D2, and the generator shifts its density from the initial synthesized examples towards the revised synthesized examples in Step G2 by reconstructing the latter.

In terms of energy function, the descriptor shifts its low energy regions towards the observed examples, and it induces the generator to map the latent factors to its low energy regions. It achieves that by stochastically relaxing the synthesized examples towards low energy regions, and let the generator track the synthesized examples.

True to the essence of learning, the two nets still learn from the data. The descriptor learns from the observed data and teaches the generator through the synthesized data. Specifically, they collaborate and communicate with each other through synthesized data in the following three versions (S stands for synthesis). (S1) The generator generates the initial draft. (S2) The descriptor revises the draft. (S3) The generator reconstructs the revised draft.

## 2. Related work

Previous authors have already told fascinating tales of two nets, although the plot of our tale is among the more intricate.

The most fabled tale of two nets is about the generative adversarial networks (GAN) (Goodfellow et al., 2014; Denton et al., 2015; Radford et al., 2015). In GAN, the generator net is paired with a discriminator net. The two nets play adversarial roles. In our work, the generator net and the descriptor net play cooperative roles, and they feed each other the synthesized data. The learning of both nets is based on maximum likelihood, and the learning process is quite stable because of the cooperative nature and the consistent direction of the two maximum likelihood training algorithms.

The connection between the descriptor and the discriminator has been explored by (Xie et al., 2016), where the descriptor can be derived from the discriminator. See also earlier pioneering work by (Tu, 2007; Welling et al., 2002).

Our tale of two nets is most similar to the one in the recent pioneering work of (Kim & Bengio, 2016). In fact, the settings of the two nets are exactly the same. In their work, the generator is learned from the current descriptor by minimizing the Kullback-Leibler divergence from the generator to the current descriptor, which can be decomposed into an energy term and an entropy term, with the latter approximated based on quantities produced by batch normalization (Ioffe & Szegedy, 2015). Their work avoids MCMC sampling, in particular, their work does not involve revising and reconstructing the synthesized examples, which plays a pivotal role in our method. In their work, the descriptor acts on the generator directly, whereas in our work, the descriptor acts on the generator through the synthesized data.

Despite the difference, there is an interesting connection between our work and (Kim & Bengio, 2016): while (Kim & Bengio, 2016) seeks to minimize the Kullback-Leibler divergence from the generator to the current descriptor, the Langevin revision dynamics in our work runs a Markov chain from the current generator to the current descriptor, and it automatically and monotonically decrease the Kullback-Leibler divergence from the distribution of the revised synthesized data to the distribution of the current descriptor (to zero in the time limit), which is an important property of the Markov chain that underlies the second law of thermodynamics and the so-called "arrow of time." (Cover & Thomas, 2012) The revised synthesized data are then fed to the generator for it to reconstruct the effect of the Langevin revisions. Therefore our learning method is consistent with that of (Kim & Bengio, 2016). In fact, the energy and entropy terms in the learning objective function in (Kim & Bengio, 2016) correspond to the gradient and noise terms of the Langevin dynamics in our work, but the Langevin dynamics relieves us the need to approximate the intractable entropy term.

Our work is related to the contrastive divergence algorithm (Hinton, 2002) for training the descriptor net. The contrastive divergence initializes the MCMC sampling from

the observed examples. The CoopNets algorithm initializes the Langevin sampling from the examples supplied by the generator.

Our work is somewhat related to the generative stochastic network of (Thibodeau-Laufer et al., 2014), which learns the Markov transition probability directly. In our work, the generator serves to absorb the cumulative effect of all the past Markov transitions (i.e., the Langevin revisions) of the descriptor and reproduce it in one shot.

Our work appears to be related to knowledge distilling (Hinton et al., 2015). In our work, the descriptor distills its knowledge to the generator.

Our work bears similarity to the co-training method of (Blum & Mitchell, 1998). The two learners in our work are of different directions, and they feed each other the synthesized data, instead of class labels.

## 3. Two nets and two algorithms

This section reviews the descriptor net and the generator net. The review of the descriptor net mainly follows (Xie et al., 2016). The review of the generator net mainly follows (Han et al., 2016).

### 3.1. Descriptor net and training algorithm

The descriptor net has its root in exponential family model in statistics and in energy-based model in machine learning. Let $Y$ be the signal, such as an image, a video sequence, or a sound sequence. The descriptor model is in the form of exponential tilting of a reference distribution (Xie et al., 2016):

$$P_{\mathcal{D}}(Y; W_{\mathcal{D}}) = \frac{1}{Z(W_{\mathcal{D}})} \exp\left[f(Y; W_{\mathcal{D}})\right] q(Y). \tag{2}$$

$q(Y)$ is the reference distribution such as Gaussian white noise

$$q(Y) = \frac{1}{(2\pi s^2)^{D/2}} \exp\left[-\frac{\|Y\|^2}{2s^2}\right], \tag{3}$$

where $D$ is the dimensionality of the signal $Y$, and $Y \sim \mathrm{N}(0, s^2 I_D)$ under $q(Y)$ ($I_D$ denotes the $D$-dimensional identity matrix). $f(Y; W_{\mathcal{D}})$ ($f$ stands for features) is the feature statistics or energy function, defined by a ConvNet whose parameters are denoted by $W_{\mathcal{D}}$. This ConvNet is bottom-up because it maps the signal $Y$ to the feature statistics. See the diagram in (1). $Z(W_{\mathcal{D}}) = \int \exp\left[f(Y; W_{\mathcal{D}})\right] q(Y) dY = \mathrm{E}_q\{\exp[f(Y; W_{\mathcal{D}})]\}$ is the normalizing constant, where $\mathrm{E}_q$ is the expectation with respect to $q$.

Suppose we observe training examples $\{Y_i, i = 1, ..., n\}$ from an unknown data distribution $P_{\mathrm{data}}(Y)$. The maximum likelihood training seeks to maximize the log-likelihood function

$$L_{\mathcal{D}}(W_{\mathcal{D}}) = \frac{1}{n} \sum_{i=1}^{n} \log P_{\mathcal{D}}(Y_i; W_{\mathcal{D}}). \tag{4}$$

If the sample size $n$ is large, the maximum likelihood estimator minimizes $\mathrm{KL}(P_{\mathrm{data}} | P_{\mathcal{D}})$, the Kullback-Leibler divergence from the data distribution $P_{\mathrm{data}}$ to the model distribution $P_{\mathcal{D}}$.

The gradient of the $L_{\mathcal{D}}(W_{\mathcal{D}})$ is

$$\begin{aligned} L'_{\mathcal{D}}(W_{\mathcal{D}}) &= \frac{1}{n} \sum_{i=1}^{n} \frac{\partial}{\partial W_{\mathcal{D}}} f(Y_i; W_{\mathcal{D}}) \\ &- \mathrm{E}_{W_{\mathcal{D}}}\left[\frac{\partial}{\partial W_{\mathcal{D}}} f(Y; W_{\mathcal{D}})\right], \end{aligned} \tag{5}$$

where $\mathrm{E}_{W_{\mathcal{D}}}$ denotes the expectation with respect to $P_{\mathcal{D}}(Y; W_{\mathcal{D}})$. The key to the above identity is that $\frac{\partial}{\partial W_{\mathcal{D}}} \log Z(W_{\mathcal{D}}) = \mathrm{E}_{W_{\mathcal{D}}}[\frac{\partial}{\partial W_{\mathcal{D}}} f(Y; W_{\mathcal{D}})]$.

The expectation in equation (5) is analytically intractable and has to be approximated by MCMC, such as Langevin revision dynamics, which iterates the following steps:

$$Y_{\tau+1} = Y_\tau - \frac{\delta^2}{2}\left[\frac{Y_\tau}{s^2} - \frac{\partial}{\partial Y} f(Y_\tau; W_{\mathcal{D}})\right] + \delta U_\tau, \tag{6}$$

where $\tau$ indexes the time steps of the Langevin dynamics, $\delta$ is the step size, and $U_\tau \sim \mathrm{N}(0, I_D)$ is the Gaussian white noise term.

We can run $\tilde{n}$ parallel chains of Langevin dynamics according to (6) to obtain the synthesized examples $\{\tilde{Y}_i, i = 1, ..., \tilde{n}\}$. The Monte Carlo approximation to $L'_{\mathcal{D}}(W_{\mathcal{D}})$ is

$$\begin{aligned} L'_{\mathcal{D}}(W_{\mathcal{D}}) &\approx \frac{1}{n} \sum_{i=1}^{n} \frac{\partial}{\partial W_{\mathcal{D}}} f(Y_i; W_{\mathcal{D}}) \\ &- \frac{1}{\tilde{n}} \sum_{i=1}^{\tilde{n}} \frac{\partial}{\partial W_{\mathcal{D}}} f(\tilde{Y}_i; W_{\mathcal{D}}), \end{aligned} \tag{7}$$

which is the difference between the observed examples and the synthesized examples. See (Lu et al., 2016) for details.

Algorithm 1 (Xie et al., 2016) describes the training algorithm. See Figure 1 for an illustration. Step D1 needs to compute $\frac{\partial}{\partial Y} f(Y; W_{\mathcal{D}})$. Step D2 needs to compute $\frac{\partial}{\partial W_{\mathcal{D}}} f(Y; W_{\mathcal{D}})$. The computations of both derivatives can be powered by back-propagation. Because of the ConvNet structure of $f(Y; W_{\mathcal{D}})$, the computations of the two derivatives share most of their steps in the chain rule computations.

Because the parameter $W_{\mathcal{D}}$ keeps changing in the learning process, the energy landscape and the local energy minima

**Algorithm 1** Algorithm D

**Input:**

(1) training examples $\{Y_i, i = 1, ..., n\}$

(2) number of Langevin steps $l_{\mathcal{D}}$

(3) number of learning iterations $T$

**Output:**

(1) estimated parameters $W_{\mathcal{D}}$

(2) synthesized examples $\{\tilde{Y}_i, i = 1, ..., \tilde{n}\}$

1: Let $t \leftarrow 0$, initialize $W_{\mathcal{D}}$.
2: Initialize $\tilde{Y}_i$, $i = 1, ..., \tilde{n}$.
3: **repeat**
4: **Step D1 Langevin revision**: For each $i$, run $l_{\mathcal{D}}$ steps of Langevin dynamics to update $\tilde{Y}_i$, i.e., starting from the current $\tilde{Y}_i$, each step follows equation (6).
5: **Step D2 density shifting**: Update $W_{\mathcal{D}}^{(t+1)} = W_{\mathcal{D}}^{(t)} + \gamma_t L'_{\mathcal{D}}(W_{\mathcal{D}}^{(t)})$, with learning rate $\gamma_t$, where $L'_{\mathcal{D}}(W_{\mathcal{D}}^{(t)})$ is computed according to (7).
6: Let $t \leftarrow t + 1$
7: **until** $t = T$

also keep changing. This may help the Langevin revision dynamics avoid being trapped by the local energy minima.

Algorithm D is a stochastic approximation algorithm (Robbins & Monro, 1951), except that the synthesized examples are obtained by a finite number of Langevin steps in each learning iteration. The convergence of an algorithm of this type to the maximum likelihood estimate has been established by (Younes, 1999).

### 3.2. Generator net and training algorithm

The generator net (Goodfellow et al., 2014) has its root in factor analysis in statistics and in latent variable model or directed graphical model in machine learning. The generator net seeks to explain the signal $Y$ of dimension $D$ by a vector of latent factors $X$ of dimension $d$, and usually $d \ll D$. The model is of the following form:

$$\begin{aligned} &X \sim \mathrm{N}(0, I_d), \\ &Y = g(X; W_{\mathcal{G}}) + \epsilon, \ \epsilon \sim \mathrm{N}(0, \sigma^2 I_D). \end{aligned} \tag{8}$$

$g(X; W_{\mathcal{G}})$ ($g$ stands for generator) is a top-down ConvNet defined by the parameters $W_{\mathcal{G}}$. The ConvNet $g$ maps the latent factors $X$ to the signal $Y$. See the diagram in (1).

The joint density of model (8) is $P_{\mathcal{G}}(X, Y; W_{\mathcal{G}}) = P_{\mathcal{G}}(X) P_{\mathcal{G}}(Y|X; W_{\mathcal{G}})$, and

$$\begin{aligned} \log P_{\mathcal{G}}(X, Y; W_{\mathcal{G}}) &= -\frac{1}{2\sigma^2} \|Y - g(X; W_{\mathcal{G}})\|^2 \\ &- \frac{1}{2}\|X\|^2 + \text{constant}, \end{aligned} \tag{9}$$

where the constant term is independent of $X$, $Y$ and $W_{\mathcal{G}}$. The marginal density is obtained by integrating out the latent factors $X$, i.e., $P_{\mathcal{G}}(Y; W_{\mathcal{G}}) = \int P_{\mathcal{G}}(X, Y; W_{\mathcal{G}}) dX$. The inference of $X$ given $Y$ is based on the posterior density $P_{\mathcal{G}}(X|Y; W_{\mathcal{G}}) = P_{\mathcal{G}}(X, Y; W_{\mathcal{G}})/P_{\mathcal{G}}(Y; W_{\mathcal{G}}) \propto P_{\mathcal{G}}(X, Y; W_{\mathcal{G}})$ as a function of $X$.

For the training data $\{Y_i, i = 1, ..., n\}$, the generator net can be trained by maximizing the log-likelihood

$$L_{\mathcal{G}}(W_{\mathcal{G}}) = \frac{1}{n} \sum_{i=1}^{n} \log P_{\mathcal{G}}(Y_i; W_{\mathcal{G}}). \tag{10}$$

For large sample, $W_{\mathcal{G}}$ is obtained by minimizing the Kullback-Leibler divergence $\mathrm{KL}(P_{\text{data}}|P_{\mathcal{G}})$ from the data distribution $P_{\text{data}}$ to the model distribution $P_{\mathcal{G}}$.

The gradient of $L_{\mathcal{G}}(W_{\mathcal{G}})$ is obtained according to the following identity

$$\begin{aligned} &\frac{\partial}{\partial W_{\mathcal{G}}} \log P_{\mathcal{G}}(Y; W_{\mathcal{G}}) \\ = \ &\mathrm{E}_{P_{\mathcal{G}}(X|Y;W_{\mathcal{G}})} \left[ \frac{\partial}{\partial W_{\mathcal{G}}} \log P_{\mathcal{G}}(X, Y; W_{\mathcal{G}}) \right]. \end{aligned} \tag{11}$$

The above identity underlies the EM algorithm. The usefulness of identify (11) lies in the fact that the derivative of the complete-data log-likelihood $\log P_{\mathcal{G}}(X, Y; W_{\mathcal{G}})$ on the right hand side can be obtained in closed form.

In general, the expectation in (11) is analytically intractable, and has to be approximated by MCMC that samples from the posterior $P_G(X|Y; W_{\mathcal{G}})$, such as the Langevin inference dynamics, which iterates

$$X_{\tau+1} = X_\tau + \frac{\delta^2}{2} \frac{\partial}{\partial X} \log P_{\mathcal{G}}(X_\tau, Y; W_{\mathcal{G}}) + \delta U_\tau, \tag{12}$$

where $\tau$ indexes the time step, $\delta$ is the step size, and for notational simplicity, we continue to use $U_\tau$ to denote the noise term, but here $U_\tau \sim \mathrm{N}(0, I_d)$. The Langevin inference solves a $\ell_2$ penalized non-linear least squares problem so that $X_i$ can reconstruct $Y_i$ given the current $W_{\mathcal{G}}$. The Langevin inference process performs explaining-away reasoning, where the latent factors in $X$ compete with each other to explain the current residual $Y - g(X; W_{\mathcal{G}})$.

With $X_i$ sampled from $P_{\mathcal{G}}(X_i \mid Y_i, W_{\mathcal{G}})$ for each observation $Y_i$ by the Langevin inference process, the Monte Carlo approximation to $L'_{\mathcal{G}}(W_{\mathcal{G}})$ is

$$\begin{aligned} L'_{\mathcal{G}}(W_{\mathcal{G}}) &\approx \frac{1}{n} \sum_{i=1}^{n} \frac{\partial}{\partial W_{\mathcal{G}}} \log P_{\mathcal{G}}(X_i, Y_i; W_{\mathcal{G}}) \\ &= \frac{1}{n} \sum_{i=1}^{n} \frac{1}{\sigma^2} (Y_i - g(X_i; W_{\mathcal{G}})) \frac{\partial}{\partial W_{\mathcal{G}}} g(X_i; W_{\mathcal{G}}). \end{aligned} \tag{13}$$

The updating of $W_{\mathcal{G}}$ solves a non-linear regression problem, so that the learned $W_{\mathcal{G}}$ enables better reconstruction

of $Y_i$ by the inferred $X_i$. Given the inferred $X_i$, the learning of $W_{\mathcal{G}}$ is a supervised learning problem (Dosovitskiy et al., 2015).

**Algorithm 2** Algorithm G

**Input:**

(1) training examples $\{Y_i, i = 1, ..., n\}$
(2) number of Langevin steps $l_{\mathcal{G}}$
(3) number of learning iterations $T$

**Output:**

(1) estimated parameters $W_{\mathcal{G}}$
(2) inferred latent factors $\{X_i, i = 1, ..., n\}$

1: Let $t \leftarrow 0$, initialize $W_{\mathcal{G}}$.
2: Initialize $X_i$, $i = 1, ..., n$.
3: **repeat**
4: **Step G1 Langevin inference**: For each $i$, run $l_{\mathcal{G}}$ steps of Langevin dynamics to update $X_i$, i.e., starting from the current $X_i$, each step follows equation (12).
5: **Step G2 reconstruction**: Update $W_{\mathcal{G}}^{(t+1)} = W_{\mathcal{G}}^{(t)} + \gamma_t L_{\mathcal{G}}{}'(W_{\mathcal{G}}^{(t)})$, with learning rate $\gamma_t$, where $L_{\mathcal{G}}{}'(W_{\mathcal{G}}^{(t)})$ is computed according to equation (13).
6: Let $t \leftarrow t + 1$
7: **until** $t = T$

Algorithm 2 (Han et al., 2016) describes the training algorithm. See Figure 2 for an illustration. Step G1 needs to compute $\frac{\partial}{\partial X} g(X; W_{\mathcal{G}})$. Step G2 needs to compute $\frac{\partial}{\partial W_{\mathcal{G}}} g(X; W_{\mathcal{G}})$. The computations of both derivatives can be powered by back-propagation, and the computations of the two derivatives share most of their steps in the chain rule computations.

Algorithm G is a stochastic approximation or stochastic gradient algorithm that converges to the maximum likelihood estimate (Younes, 1999) .

## 4. CoopNets algorithm

In Algorithm D and Algorithm G, both steps D1 and G1 are Langevin dynamics. They may be slow to converge and may become bottlenecks in their respective algorithms. An interesting observation is that the two algorithms can cooperate with each other by jumpstarting each other's Langevin sampling.

Specifically, in Step D1, we can initialize the synthesized examples by generating examples from the generator net, which does not require MCMC, because the generator net is a directed graphical model. More specifically, we first generate $\hat{X}_i \sim \mathrm{N}(0, I_d)$, and then generate $\hat{Y}_i = g(\hat{X}_i; W_{\mathcal{G}}) + \epsilon_i$, for $i = 1, ..., \tilde{n}$. If the current generator $P_{\mathcal{G}}$ is close to the current descriptor $P_{\mathcal{D}}$, then the generated $\{\hat{Y}_i\}$ should be a good initialization for sampling from the descriptor net, i.e., starting from the $\{\hat{Y}_i, i = 1, ..., \tilde{n}\}$ supplied by the generator net, we run Langevin dynamics in Step D1 for $l_{\mathcal{D}}$ steps to get $\{\tilde{Y}_i, i = 1, ..., \tilde{n}\}$, which are revised versions of $\{\hat{Y}_i\}$. These $\{\tilde{Y}_i\}$ can be used as the synthesized examples from the descriptor net. We can then update $W_{\mathcal{D}}$ according to Step D2 of Algorithm D.

In order to update $W_{\mathcal{G}}$ of the generator net, we treat the $\{\tilde{Y}_i, i = 1, ..., \tilde{n}\}$ produced by the above D1 step as the training data for the generator. Since these $\{\tilde{Y}_i\}$ are obtained by the Langevin revision dynamics initialized from the $\{\hat{Y}_i, i = 1, ..., \tilde{n}\}$ produced by the generator net with known latent factors $\{\hat{X}_i, i = 1, ..., \tilde{n}\}$, we can update $W_{\mathcal{G}}$ by learning from $\{(\tilde{Y}_i, \hat{X}_i), i = 1, ..., \tilde{n}\}$, which is a supervised learning problem, or more specifically, a non-linear regression of $\tilde{Y}_i$ on $\hat{X}_i$. At $W_{\mathcal{G}}^{(t)}$, the latent factors $\hat{X}_i$ generates and thus reconstructs the initial example $\hat{Y}_i$. After updating $W_{\mathcal{G}}$, we want $\hat{X}_i$ to reconstruct the revised example $\tilde{Y}_i$. That is, we revise $W_{\mathcal{G}}$ to absorb the revision from $\hat{Y}_i$ to $\tilde{Y}_i$, so that the generator shifts its density from $\{\hat{Y}_i\}$ to $\{\tilde{Y}_i\}$. The reconstruction error can tell us whether the generator has caught up with the descriptor by fully absorbing the revision.

The left diagram in (14) illustrates the basic idea.

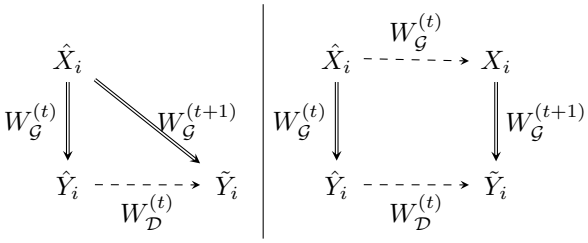


(14)

In the two diagrams in (14), the double line arrows indicate generation and reconstruction in the generator net, while the dashed line arrows indicate Langevin dynamics for revision and inference in the two nets. The diagram on the right in (14) illustrates a more rigorous method, where we initialize the Langevin inference of $\{X_i, i = 1, ..., \tilde{n}\}$ in Step G1 from $\{\hat{X}_i\}$, and then update $W_{\mathcal{G}}$ in Step G2 based on $\{(\tilde{Y}_i, X_i), i = 1, ..., \tilde{n}\}$.

Algorithm 3 describes the cooperative training that interweaves Algorithm D and Algorithm G. See Figure 3 for the flow chart of the CoopNets algorithm. In our experiments, we set $l_{\mathcal{G}} = 0$ and infer $X_i = \hat{X}_i$.

The learning of both the descriptor and the generator follows the "analysis by synthesis" principle (Grenander & Miller, 2007). There are three sets of synthesized examples (S stands for synthesis). (S1) Initial synthesized examples $\{\hat{Y}_i\}$ generated by Step G0. (S2) Revised synthesized examples $\{\tilde{Y}_i\}$ produced by Step D1. (S3) Reconstructed synthesized examples $\{g(X_i; W_{\mathcal{G}}^{(t+1)})\}$ produced by Step

**Algorithm 3** CoopNets Algorithm

**Input:**

(1) training examples $\{Y_i, i = 1, ..., n\}$

(2) numbers of Langevin steps $l_{\mathcal{D}}$ ad $l_{\mathcal{G}}$

(3) number of learning iterations $T$

**Output:**

(1) estimated parameters $W_{\mathcal{D}}$ and $W_{\mathcal{G}}$

(2) synthesized examples $\{\tilde{Y}_i, \tilde{Y}_i, i = 1, ..., \tilde{n}\}$

1: Let $t \leftarrow 0$, initialize $W_{\mathcal{D}}$ and $W_{\mathcal{G}}$.

2: **repeat**

3: **Step G0**: For $i = 1, ..., \tilde{n}$, generate $\hat{X}_i \sim \mathrm{N}(0, I_d)$, and generate $\hat{Y}_i = g(\hat{X}_i; W_{\mathcal{G}}^{(t)}) + \epsilon_i$.

4: **Step D1**: For $i = 1, ..., \tilde{n}$, starting from $\hat{Y}_i$, Run $l_{\mathcal{D}}$ steps of Langevin revision dynamics to obtain $\tilde{Y}_i$, each step following equation (6).

5: **Step G1**: Treat the current $\{\tilde{Y}_i, i = 1, ..., \tilde{n}\}$ as the training data, for each $i$, infer $X_i = \hat{X}_i$. Or more rigorously, starting from $X_i = \hat{X}_i$, run $l_{\mathcal{G}}$ steps of Langevin inference dynamics to update $X_i$, each step following equation (12).

6: **Step D2**: Update $W_{\mathcal{D}}^{(t+1)} = W_{\mathcal{D}}^{(t)} + \gamma_t L'_{\mathcal{D}}(W_{\mathcal{D}}^{(t)})$, where $L'_{\mathcal{D}}(W_{\mathcal{D}}^{(t)})$ is computed according to (7).

7: **Step G2**: Update $W_{\mathcal{G}}^{(t+1)} = W_{\mathcal{G}}^{(t)} + \gamma_t {L_{\mathcal{G}}}'(W_{\mathcal{G}}^{(t)})$, where ${L_{\mathcal{G}}}'(W_G)$ is computed according to equation (13), except that $Y_i$ is replaced by $\tilde{Y}_i$, and $n$ by $\tilde{n}$. We can run multiple iterations of Step G2 to learn from and reconstruct $\{\tilde{Y}_i\}$, and to allow the generator to catch up with the descriptor.

8: Let $t \leftarrow t + 1$

9: **until** $t = T$

G2. The descriptor shifts its density from (S2) towards the observed data, while the generator shifts its density from (S1) towards (S2) (by arriving at (S3)).

The evolution from (S1) to (S2) is the work of the descriptor. It is a process of stochastic relaxation that settles the synthesized examples in the low energy regions of the descriptor. The descriptor works as an associative memory, with (S1) being the cue, and (S2) being the recalled memory. It serves as a feedback to the generator. The reconstruction of (S2) by (S3) is the work of the generator that seeks to absorb the feedback conveyed by the evolution from (S1) to (S2). The descriptor can test whether the generator learns well by checking whether (S3) is close to (S2). The two nets collaborate and communicate with each other via synthesized data.

## 5. Experiments

We use the MatConvNet of (Vedaldi & Lenc, 2015) for coding. For the descriptor net, we adopt the structure of (Xie et al., 2016), where the bottom-up network consists of multiple layers of convolution by linear filtering, ReLU non-linearity, and down-sampling. We adopt the structure of the generator network of (Radford et al., 2015; Dosovitskiy et al., 2015), where the top-down network consists of multiple layers of deconvolution by linear superposition, ReLU non-linearity, and up-sampling, with tanh non-linearity at the bottom-layer (Radford et al., 2015) to make the signals fall within $[-1, 1]$. The synthesized images we show are all generated by the descriptor net.

### 5.1. Learning texture patterns

We learn a separate model from each texture image. The training images are collected from the Internet, and then resized to 224 $\times$ 224. The synthesized images are of the same size as the training images.

We use a 3-layer descriptor net, where the first layer has 100 15 $\times$ 15 filters with sub-sampling rate of 3 pixels, the second layer has 70 9 $\times$ 9 filters with sub-sampling of 1, and the third layer has 30 7 $\times$ 7 filters with sub-sampling of 1. We fix the standard deviation of the reference distribution of the descriptor net to be $s = 0.012$. We use $L$=20 or 30 steps of Langevin revision dynamics within each learning iteration, and the Langevin step size is set at 0.003. The learning rate is 0.01.

Starting from 7 $\times$ 7 latent factors, the generator net has 5 layers of deconvolution with 5 $\times$ 5 kernels (basis functions), with an up-sampling factor of 2 at each layer. The standard deviation of the noise vector is $\sigma = 0.3$. The learning rate is $10^{-6}$. The number of generator learning steps is 1 at each cooperative learning iteration.

We run $10^4$ cooperative learning iterations to train the CoopNets. Figure 4 displays the results of generating texture patterns. For each category, the first image is the training image, and the rest are 3 of the images generated by the learning algorithm. We run $\tilde{n} = 6$ parallel chains for the first example, where images from 3 of them are presented. We run a single chain for the rest examples, where the synthesized images are generated at different iterations. Even though we run a single chain, it is as if we run an infinite number of chains, because in each iteration, we run Langevin revision dynamics from a new image sampled from the generator.

### 5.2. Learning object patterns

We learn a separate model for each object pattern. The training images are collected from the Internet, and then resized to 64 $\times$ 64. The number of training images for each category is about 5. The structure of the generator network is the same as (Radford et al., 2015; Dosovitskiy et al., 2015). We adopt a 3-layer descriptor net. The first

*Figure 4.* Generating texture patterns. Each row displays one texture experiment, where the first image is the training image, and the rest are 3 of the images generated by the CoopNets algorithm.

*Figure 5.* Generating object patterns. Each row displays one object experiment, where the first 4 images are 4 of the training images, and the rest are 4 of the images generated by the CoopNets algorithm.

*Figure 6.* Generating human face pattern. The synthesized images are generated by the coopNets algorithm that learns from $10,000$ images.

layer has 100 $4 \times 4$ filters with sub-sampling of 2, the second layers has 64 $2 \times 2$ filters with sub-sampling of 1, and the third layer is a fully connected layer with single channel output. We fix the standard deviation of the reference distribution of the descriptor net to be $s = 0.016$. We use $L$=10 steps of Langevin revision dynamics within each learning iteration, and the Langevin step size is set at 0.005. The learning rate is 0.07. We run $3,000$ cooperative learning iterations to train the model. Figure 5 displays the results. Each row displays an experiment, where the first 4 images are 4 of the training images and the rest are generated images.

We also conduct an experiment on human faces, where we adopt a 4-layer descriptor net. The first layer has 96 $5 \times 5$ filters with sub-sampling of 2, the second layers has 128 $5 \times 5$ filters with sub-sampling of 2, the third layer has 256 $5 \times 5$ filters with sub-sampling of 2, and the final layer is a fully connected layer with 50 channels as output. We reduce the Langevin step size to 0.002. The training data are $10,000$ human faces randomly selected from CelebA dataset (Liu et al., 2015). We run 600 cooperative learning iterations. Figure 6 displays 144 synthesized human faces by the descriptor net.

## 6. Conclusion

The most unique feature of our work is that the two networks feed each other the synthesized data in the learning process. The generator feeds the descriptor the initial version of the synthesized data. The descriptor feedbacks the generator the revised version of the synthesized data. The generator then produces the reconstructed version of the synthesized data. While the descriptor learns from finite amount of observed data, the generator learns from virtually infinite amount of synthesized data.

Another unique feature of our work is that the learning process interweaves the existing maximum likelihood learning algorithms for the two networks, so that the two algorithms jumpstart each other's MCMC sampling.

A third unique feature of our work is that the MCMC for the descriptor keeps rejuvenating the chains by refreshing the samples by independent replacements supplied by the generator, so that a single chain effectively amounts to an infinite number of chains or the evolution of the whole marginal distribution. Powering the MCMC sampling of the descriptive models in (Lu et al., 2016; Xie et al., 2016) is the main motivation of this paper, with the bonus of turning the unsupervised learning of the generator (Han et al., 2016) into supervised learning.

## Code and data

`http://www.stat.ucla.edu/~ywu/CoopNets/main.html`

## Acknowledgement

We thank Hansheng Jiang for her work on this project as a summer visiting student. We thank Tian Han for sharing the code on learning the generator network, and for helpful discussions.

The work is supported by NSF DMS 1310391, DARPA SIMPLEX N66001-15-C-4035, ONR MURI N00014-16-1-2007, and DARPA ARO W911NF-16-1-0579.

## 7. Appendix 1: Convergence

### 7.1. Generator of infinite capacity

In the CoopNets algorithm, the descriptor learns from the observed examples, while the generator learns from the descriptor through the synthesized examples. Therefore, the descriptor is the driving force in terms of learning, although the generator is the driving force in terms of synthesis. In order to understand the convergence of learning, we can start from Algorithm D for learning the descriptor.

Algorithm D is a stochastic approximation algorithm (Robbins & Monro, 1951), except that the samples are generated by finite step MCMC transitions. According to (Younes, 1999), Algorithm D converges to the maximum likelihood estimate under suitable regularity conditions on the mixing of the transition kernel of the MCMC and the schedule of

the learning rate $\gamma_t$, even if the number of Langevin steps $l_{\mathcal{D}}$ is finite or small (e.g., $l_{\mathcal{D}} = 1$), and even if the number of parallel chains $\tilde{n}$ is finite or small (e.g., $\tilde{n} = 1$). The reason is that the random fluctuations caused by the finite number of chains, $\tilde{n}$, and the limited mixing caused by the finite steps of MCMC, $l_{\mathcal{D}}$, are mitigated if the learning rate $\gamma_t$ is sufficiently small. At learning iteration $t$, let $W_{\mathcal{D}}^{(t)}$ be the estimated parameter of the descriptor. Let $P_{\mathcal{D}}^{(t+1)}$ be the marginal distribution of $\{\tilde{Y}_i\}$, even though $P_{\mathcal{D}}^{(t+1)} \neq P_{\mathcal{D}}(Y; W_{\mathcal{D}}^{(t)})$ because $l_{\mathcal{D}}$ is finite ($P_{\mathcal{D}}^{(t+1)} = P_{\mathcal{D}}(Y; W_{\mathcal{D}}^{(t)})$ if $l_{\mathcal{D}} \to \infty$), we still have $W_{\mathcal{D}}^{(t)} \to \hat{W}_{\mathcal{D}}$ in probability according to (Younes, 1999), where $\hat{W}_{\mathcal{D}}$ is the maximum likelihood estimate of $W_{\mathcal{D}}$.

The efficiency of Algorithm D increases if the number of parallel chains $\tilde{n}$ is large because it leads to more accurate estimation of the expectation in the gradient $L'_{\mathcal{D}}(W_{\mathcal{D}})$ of equation (5), so that we can afford to use larger learning rate $\gamma_t$ for faster convergence.

Now let us come back to the CoopNets algorithm. In order to understand how the descriptor net helps the training of the generator net, let us consider the idealized scenario where the number of parallel chains $\tilde{n} \to \infty$, and the generator has infinite capacity, and in each iteration it estimates $W_{\mathcal{G}}$ by maximum likelihood using the synthesized data from $P_{\mathcal{D}}^{(t+1)}$. In this idealized scenario, the learned generator $P_{\mathcal{G}}(Y; W_{\mathcal{G}}^{(t+1)})$ will reproduce $P_{\mathcal{D}}^{(t+1)}$ by minimizing $\mathrm{KL}(P_{\mathcal{D}}^{(t+1)}(Y)|P_{\mathcal{G}}(Y; W_{\mathcal{G}}))$, with $P_{\mathcal{D}}^{(t+1)}$ serving as its data distribution. Then eventually the learned generator $P_{\mathcal{G}}(Y, \hat{W}_{\mathcal{G}})$ will reproduce $P_{\mathcal{D}}(Y; \hat{W}_{\mathcal{D}})$. Thus the cooperative training helps the learning of the generator. Note that the learned generator $P_{\mathcal{G}}(Y, \hat{W}_{\mathcal{G}})$ will not reproduce the distribution of the observed data $P_{\mathrm{data}}$, unless the descriptor is of infinite capacity too.

Conversely, the generator net also helps the learning of the descriptor net in the CoopNets algorithm. In Algorithm D, it is impractical to make the number of parallel chains $\tilde{n}$ too large. On the other hand, it would be difficult for a small number of chains $\{\tilde{Y}_i, i = 1, ..., \tilde{n}\}$ to explore the state space. In the CoopNets algorithm, because $P_{\mathcal{G}}(Y; W_{\mathcal{G}}^{(t)})$ reproduces $P_{\mathcal{D}}^{(t)}$, we can generate a completely new batch of independent samples $\{\hat{Y}_i\}$ from $P_{\mathcal{G}}(Y; W_{\mathcal{G}}^{(t)})$, and revise $\{\hat{Y}_i\}$ to $\{\tilde{Y}_i\}$ by Langevin dynamics, instead of running Langevin dynamics from the same old batch of $\{\tilde{Y}_i\}$ as in the original Algorithm D. This is like implementing an infinite number of parallel chains, because each iteration evolves a fresh batch of examples, as if each iteration evolves a new set of chains. By updating the generator $W_{\mathcal{G}}$, it is like we are updating the infinite number of parallel chains, because $W_{\mathcal{G}}$ memorizes the whole distribution. Even if $\tilde{n}$ in the CoopNets algorithm is small, e.g., $\tilde{n} = 1$, viewed from the perspective of Algorithm D, it is as if $\tilde{n} \to \infty$. Thus the above idealization $\tilde{n} \to \infty$ is sound.

### 7.2. Generator of finite capacity

From an information geometry (Amari & Nagaoka, 2007) point of view, let $\mathcal{D} = \{P_{\mathcal{D}}(Y; W_{\mathcal{D}}), \forall W_{\mathcal{D}}\}$ be the manifold of the descriptor models, where each distribution $P_{\mathcal{D}}(Y; W_{\mathcal{D}})$ is a point on this manifold. Then the maximum likelihood estimate of $W_{\mathcal{D}}$ is a projection of the data distribution $P_{\mathrm{data}}$ onto the manifold $\mathcal{D}$. Let $\mathcal{G} = \{P_{\mathcal{G}}(Y; W_{\mathcal{G}}), \forall W_{\mathcal{G}}\}$ be the manifold of the generator models, where each distribution $P_{\mathcal{G}}(Y; W_{\mathcal{G}})$ is a point on this manifold. Then the maximum likelihood estimate of $W_{\mathcal{G}}$ is a projection of the data distribution $P_{\mathrm{data}}$ onto the manifold $\mathcal{G}$.

From now on, for notational simplicity and with a slight abuse of notation, we use $W_{\mathcal{D}}$ to denote the descriptor distribution $P_{\mathcal{D}}(Y; W_{\mathcal{D}})$, and use $W_{\mathcal{G}}$ to denote the generator distribution $P_{\mathcal{G}}(Y; W_{\mathcal{G}})$.

We assume both the observed data size $n$ and the synthesized data size $\tilde{n}$ are large enough so that we shall work on distributions or populations instead of finite samples. As explained above, assuming $\tilde{n} \to \infty$ is sound because the generator net can supply unlimited number of examples.

The Langevin revision dynamics runs a Markov chain from $W_{\mathcal{G}}^{(t)}$ towards $W_{\mathcal{D}}^{(t)}$. Let $\mathbf{L}_{W_{\mathcal{D}}}$ be the Markov transition kernel of $l_{\mathcal{D}}$ steps of Langevin revisions towards $W_{\mathcal{D}}$. The distribution of the revised synthesized data is

$$P_{\mathcal{D}}^{(t+1)} = \mathbf{L}_{W_{\mathcal{D}}^{(t)}} \cdot W_{\mathcal{G}}^{(t)}, \tag{15}$$

where the notation $\mathbf{L} \cdot P$ denotes the marginal distribution obtained by running the Markov transition $\mathbf{L}$ from $P$. The distribution $P_{\mathcal{D}}^{(t+1)}$ is in the middle between the two nets $W_{\mathcal{G}}^{(t)}$ and $W_{\mathcal{D}}^{(t)}$, and it serves as the data distribution to train the generator, i.e., we project this distribution to the manifold $\mathcal{G} = \{P_{\mathcal{G}}(Y; W_{\mathcal{G}}), \forall W_{\mathcal{G}}\} = \{W_{\mathcal{G}}\}$ (recall we use $W_{\mathcal{G}}$ to denote the distribution $P_{\mathcal{G}}(Y; W_{\mathcal{G}})$) in the information geometry picture, so that

$$W_{\mathcal{G}}^{(t+1)} = \arg\min_{\mathcal{G}} \mathrm{KL}(P_{\mathcal{D}}^{(t+1)}|W_{\mathcal{G}}). \tag{16}$$

The learning process alternates between Markov transition in (15) and projection in (16), as illustrated by Figure 7. Compared to the three sets of synthesized data mentioned at the end of Section 4, $W_{\mathcal{G}}^{(t)}$ corresponds to (S1), $P_{\mathcal{D}}^{(t+1)}$ corresponds to (S2), and $W_{\mathcal{G}}^{(t+1)}$ corresponds to (S3).

In the case of $l_{\mathcal{D}} \to \infty$,

$$W_{\mathcal{D}}^{(t)} \to \hat{W}_{\mathcal{D}} = \arg\min_{\mathcal{D}} \mathrm{KL}(P_{\mathrm{data}}|W_{\mathcal{D}}), \tag{17}$$

$$W_{\mathcal{G}}^{(t)} \to \hat{W}_{\mathcal{G}} = \arg\min_{\mathcal{G}} \mathrm{KL}(\hat{W}_{\mathcal{D}}|W_{\mathcal{G}}). \tag{18}$$

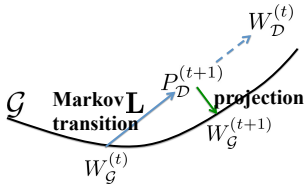


*Figure 7.* The learning of the generator alternates between Markov transition and projection. The family of the generator models $\mathcal{G}$ is illustrated by the black curve. Each distribution is illustrated by a point.

That is, we first project $P_{\text{data}}$ onto $\mathcal{D}$, and from there continue to project onto $\mathcal{G}$. Therefore, $W_{\mathcal{D}}$ converges to the maximum likelihood estimate with $P_{\text{data}}$ being the data distribution, and $W_{\mathcal{G}}$ converges to the maximum likelihood estimate with $\hat{W}_{\mathcal{D}}$ serving as the data distribution.

For finite $l_{\mathcal{D}}$, the algorithm may converge to the following fixed points. The fixed point for the generator satisfies

$$\hat{W}_{\mathcal{G}} = \arg\min_{\mathcal{G}} \text{KL}(L_{\hat{W}_{\mathcal{D}}} \cdot \hat{W}_{\mathcal{G}} | W_{\mathcal{G}}). \tag{19}$$

The fixed point for the descriptor satisfies

$$\begin{aligned} \hat{W}_{\mathcal{D}} \quad = \quad & \arg\min_{\mathcal{D}} \big[\text{KL}(P_{\text{data}}|W_{\mathcal{D}}) \\ & -\text{KL}(\mathbf{L}_{\hat{W}_{\mathcal{D}}} \cdot \hat{W}_{\mathcal{G}}|W_{\mathcal{D}})\big], \end{aligned} \tag{20}$$

which is similar to contrastive divergence (Hinton, 2002), except that $\hat{W}_G$ takes the place of $P_{\text{data}}$ in the second Kullback-Leibler divergence. Because $\hat{W}_{\mathcal{G}}$ is supposed to be close to $\hat{W}_{\mathcal{D}}$, the second Kullback-Leibler divergence is supposed to be small, hence our algorithm is closer to maximum likelihood learning than contrastive divergence.

(Kim & Bengio, 2016) learned the generator by gradient descent on $\text{KL}(W_{\mathcal{G}}|W_{\mathcal{D}}^{(t)})$ over $\mathcal{G}$. The objective function is $\text{KL}(W_{\mathcal{G}}|W_{\mathcal{D}}^{(t)}) = \text{E}_{W_{\mathcal{G}}}[\log P_G(Y; W_{\mathcal{G}})] - \text{E}_{W_{\mathcal{G}}}[\log P_{\mathcal{D}}(Y; W_{\mathcal{D}}^{(t)})]$, where the first term is the negative entropy that is intractable, and the second term is the expected energy that is tractable. Our learning method for the generator is consistent with the learning objective $\text{KL}(W_{\mathcal{G}}|W_{\mathcal{D}}^{(t)})$, because

$$\text{KL}(P_{\mathcal{D}}^{(t+1)}|W_{\mathcal{D}}^{(t)}) \leq \text{KL}(W_{\mathcal{G}}^{(t)}|W_{\mathcal{D}}^{(t)}). \tag{21}$$

In fact, $\text{KL}(P_{\mathcal{D}}^{(t+1)}|W_{\mathcal{D}}^{(t)}) \to 0$ monotonically as $l_{\mathcal{D}} \to \infty$ due to the second law of thermodynamics (Cover & Thomas, 2012). The reduction of the Kullback-Leibler divergence in (21) and the projection in (16) in our learning of the generator are consistent with the learning objective of reducing $\text{KL}(W_{\mathcal{G}}|W_{\mathcal{D}}^{(t)})$ in (Kim & Bengio, 2016). But the Monte Carlo implementation of $\mathbf{L}$ in our work avoids the need to approximate the intractable entropy term.

## 8. Appendix 2: Generator and MCMC

The general idea of the interaction between MCMC and the generator can be illustrated by the following diagram,

$$\begin{array}{lccc} \text{MCMC}: & P^{(t)} & \xrightarrow{\text{Markov transition}} & P^{(t+1)} \\ & \Updownarrow & & \Updownarrow \\ \text{Generator}: & W^{(t)} & \xrightarrow{\text{Parameter updating}} & W^{(t+1)} \end{array} \tag{22}$$

where $P^{(t)}$ is the marginal distribution of a MCMC and $W^{(t)}$ is the parameter of the generator which traces the evolution of the marginal distribution in MCMC by absorbing the cumulative effect of all the past Markov transitions (we drop the subscripts $\mathcal{D}$ and $\mathcal{G}$ here for simplicity and generality).

In traditional MCMC, we only have access to Monte Carlo samples, instead of their marginal distributions, which exist only theoretically but are analytically intractable. However, with the generator net, we can actually implement MCMC at the level of the whole distributions, instead of a number of Monte Carlo samples, in the sense that after learning $W^{(t)}$ from the existing samples of $P^{(t)}$, we can replace the existing samples by fresh new samples by sampling from the generator defined by the learned $W^{(t)}$, which is illustrated by the two-way arrow between $P^{(t)}$ and $W^{(t)}$. Effectively, the generator powers MCMC by implicitly running an infinite number of parallel chains. Conversely, the MCMC does not only drive the evolution of the samples, it also drives the evolution of a generator model.

If the descriptor is fixed, the generator will converge to the descriptor if the generator has infinite capacity. In the CoopNets algorithm, the descriptor is an evolving one, because we keep updating the parameter $W_{\mathcal{D}}$ of the descriptor by learning from the observed examples. Therefore the Markov chain in (22) is a non-stationary one, because the target distribution of the MCMC, i.e., the current descriptor, keeps evolving, so that the Markov transition kernel also keeps changing. Initially, the distribution of the descriptor is of high temperature, and the noise term dominates the Langevin dynamics. By following the descriptor, the generator will gain high entropy and high variance. As the learning proceeds, the parameters of the descriptor become bigger in magnitude, and the distribution becomes colder and multi-modal, so that the Langevin dynamics will be dominated by the gradient term. At the low temperature, the patterns begin to form. By chasing the descriptor, the generator begins to map the latent factors to the local modes to create the patterns. In the end, it is possible that the generator maps each small neighborhood of the factor space into a local mode of the energy landscape of the descriptor. As observed by (Xie et al., 2016), the local modes of the descriptor satisfy an auto-encoder $Y/s^2 = \partial f(Y; W_{\mathcal{D}})/\partial Y$. Hence the examples generated

by the generator $Y = g(X; W_{\mathcal{G}})$ satisfy the auto-encoder defined by the descriptor, so that the stochastic relaxation by the descriptor cannot make much revision of the examples generated by the generator. From an embedding perspective, the generator embeds the high-dimensional local modes of the descriptor into its low-dimensional factor space, thus providing a map of the energy landscape of the descriptor. From the MCMC perspective, the learning process is like a simulated annealing or tempering process (Kirkpatrick et al., 1983; Marinari & Parisi, 1992), and the generator traces out this process in distribution as illustrated by diagram (22). As is the case with simulated annealing, it is necessary to begin with a high temperature and high entropy distribution to cast a sufficiently wide generator net, and use a slow tempering schedule, or learning schedule in our case, to capture most of the major modes.

MCMC sampling of an evolving descriptor can be easier than sampling from a fixed descriptor, because the energy landscape of an evolving descriptor keeps changing or even destroying its local modes defined by the auto-encoder, thus releasing the Markov chains from the trapping of the local modes.

## 9. Appendix 3: More on two nets

### 9.1. Backgrounds

To illustrate the descriptive models and the generative models, (Zhu, 2003) used the FRAME (Filters, Random field, And Maximum Entropy) model of (Zhu et al., 1997) as a prototype of the descriptive models, and used the sparse coding model of (Olshausen & Field, 1997) as a prototype of the generative models. Both the FRAME model and the sparse coding model involve wavelets, but the wavelets play different roles in these two models. In the FRAME model, the wavelets serve as bottom-up filters for extracting features. In the sparse coding model, the wavelets serve as the top-down basis functions for linear superposition. These two roles can be unified only in an orthogonal basis or a tight frame of wavelets, but in general we cannot have both. In more general setting, in a descriptive model, the features are explicitly defined by a bottom-up process, but the signal cannot be recovered from the features by an explicit top-down process. In a generative model, the latent variables generate the signal by an explicit top-down process, but inferring the latent variables requires explaining-away reasoning that cannot be defined by an explicit bottom-up process. Efforts have been made to integrate these two classes of models and to understand their entropy regimes, such as (Guo et al., 2003; Wu et al., 2008; Xie et al., 2014; 2016), but we still lack a clear understanding of their roles, their scopes, and their relationship.

Both the descriptor net and the generator net are fundamental representations of knowledge. The descriptor provides multiple layers of features, and its proper use may be to serve as an associative memory or content addressable memory (Hopfield, 1982), where the conditional distribution of the signal given partial information is explicitly defined. Compared to the full distribution, the conditional distribution can be easier to sample from because it can be less multi-modal than the full distribution. The generator net disentangles the variations in the training examples and provides non-linear dimension reduction. It embeds the highly non-linear manifold formed by the high-dimensional training examples into the low-dimensional Euclidean space of latent factors, so that linear interpolation in the factor space results in highly non-linear interpolation in the signal space.

We could have called the descriptor net the energy-based model, but the latter term carries much broader meaning than probabilistic models, see, e.g., (LeCun et al., 2006). In fact, energy-based probabilistic models are often looked upon unfavorably because of the intractable normalizing constant and the reliance on MCMC. However, a probabilistic model appears the only coherent way to define an associative memory, where the conditional distributions given different partial informations are always consistent with each other. As to the difficulty with MCMC for sampling the descriptor net, the generator net can actually be of great help as we have explored in this paper.

There are other probabilistic models that do not belong to but are related to the above two classes of models. For example, there are energy-based models that involve latent variables, such as the restricted Boltzmann machine and its generalizations (Hinton et al., 2006; Salakhutdinov & Hinton, 2009; Lee et al., 2009). These models are undirected. There are also models where the latent variables and signals are of the same dimension, and they can be transformed to each other, so that the latent variables can also be considered features (Hyvärinen et al., 2004; Dinh et al., 2016). There are also models that impose causal auto-regressive factorizations (Oord et al., 2016; Dinh et al., 2016). These models are beyond the scope of this paper.

### 9.2. Descriptor net

The descriptor net can be considered a multi-layer generalization of the FRAME (Filters, Random field, And Maximum Entropy) model (Zhu et al., 1997; Lu et al., 2016). The energy function of the descriptor model is

$$\mathcal{E}(Y; W_{\mathcal{D}}) = \frac{\|Y\|^2}{2s^2} - f(Y; W_{\mathcal{D}}). \tag{23}$$

(Xie et al., 2016) noticed that the local energy minima satisfy an auto-encoder $\partial \mathcal{E}(Y; W_{\mathcal{D}})/\partial Y = 0$, which is

$$Y/s^2 = \partial f(Y; W_{\mathcal{D}})/\partial Y. \tag{24}$$

The right hand side involves a bottom-up pass of convolution and a top-down pass of back-propagation for computing the derivative, and this back-propagation is actually a deconvolution process (Zeiler & Fergus, 2014). This auto-encoder is curious in that back-propagation becomes representation. It also connects the Hopfield associative memory to the hierarchical distributed representation.

(Xie et al., 2016) also noticed that if the non-linearity in the bottom-up ConvNet that defines $f$ is rectified linear units (ReLU), $h(r) = \max(0, r) = 1(r > 0)r$, where $1(r > 0)$ is the binary activation, then $P_{\mathcal{D}}(Y; W_{\mathcal{D}})$ is piecewise Gaussian, and the modes of the Gaussian pieces are auto-encoding. Specifically, the activation pattern of all the ReLU units of the ConvNet partitions the space of $Y$ into different pieces, so that those $Y$ within the same piece share the same activation pattern. $f(Y; W_{\mathcal{D}})$ is linear on each piece (Pascanu et al., 2013), so that the energy function $\|Y\|^2/(2s^2) - f(Y; W_{\mathcal{D}})$ is piecewise quadratic, which gives rise to the piecewise Gaussian structure.

In the Langevin equation (6), the gradient term is in the form of the reconstruction error $Y/s^2 - \partial f(Y; W_{\mathcal{D}})/\partial Y$. If $Y$ is a local mode of the energy function, then $Y$ satisfies the auto-encoder $Y/s^2 = \partial f(Y; W_{\mathcal{D}})/\partial Y$ mentioned above. The Langevin revision dynamics explore these local modes, by stochastically relaxing the current sample towards the local mode.

(Xie et al., 2016) derived the descriptor net from the discriminator net where $f(Y; W_{\mathcal{D}})$ is the score for soft-max or multinomial logistic classification, and the reference distribution $q(Y)$ plays the role of the negative or baseline category.

### 9.3. Generator net

The generator network can be viewed a multi-layer recursion of factor analysis. (Han et al., 2016) noted that if the non-linearity in the top-down ConvNet that defines $g(X; W_{\mathcal{G}})$ is ReLU, then the model becomes piecewise linear factor analysis. Specifically, the activation pattern of all the ReLU units of the ConvNet partition the factor space of $X$ into different pieces, so that those $X$ within the same piece share the same activation pattern. Then $g(X; W_{\mathcal{G}})$ is linear on each piece (Pascanu et al., 2013), so that the model is piecewise linear factor analysis, and $P_{\mathcal{G}}(Y; W_{\mathcal{G}})$ is essentially piecewise Gaussian.

We can compare the descriptor net $P_{\mathcal{D}}(Y; W_{\mathcal{D}})$ with the generator net $P_{\mathcal{G}}(Y; W_{\mathcal{G}}) = \int P_{\mathcal{G}}(X, Y; W_{\mathcal{G}})dX$, where the integral can be approximated by the Laplace approximation around the posterior mode $\hat{X} = \hat{X}(Y; W_{\mathcal{G}})$ that maximizes $P_{\mathcal{G}}(X|Y, W_{\mathcal{G}})$. Then we may compare the energy $\|Y\|^2/(2s^2) - f(Y; W_{\mathcal{D}})$ of the descriptor with $\|Y - g(\hat{X}; W_{\mathcal{G}})\|^2/(2\sigma^2) + \|\hat{X}\|^2/2$ in the Laplace approximation of the generator. $\hat{X} = \hat{X}(Y; W_{\mathcal{G}})$ can be considered an encoder or recognizer for inferring $X$. The encoder or recognizer net has been learned in the wake-sleep algorithm (Hinton et al., 1995) and variation auto-encoder (Kingma & Welling, 2014; Rezende et al., 2014; Mnih & Gregor, 2014). Recall that $f(Y; W_{\mathcal{D}})$ can be used as score for soft-max discriminator. Thus the four nets of discriminator, descriptor, generator, and recognizer are connected.

To understand how algorithm G or its EM idealization maps the prior distribution of the latent factors $P_{\mathcal{G}}(X)$ to the data distribution $P_{\text{data}}$ by the learned $g(X; W_{\mathcal{G}})$, let us define $P_{\text{data}}(X, Y; W_{\mathcal{G}}) = P_{\text{data}}(Y)P_{\mathcal{G}}(X|Y, W_{\mathcal{G}}) = P_{\text{data}}(X; W_{\mathcal{G}})P_{\text{data}}(Y|X, W_{\mathcal{G}})$, where $P_{\text{data}}(X; W_{\mathcal{G}}) = \int P_{\mathcal{G}}(X|Y, W_{\mathcal{G}})P_{\text{data}}(Y)dY$ is obtained by averaging the posteriors $P_{\mathcal{G}}(X|Y; W_{\mathcal{G}})$ over the observed data $Y \sim P_{\text{data}}$. That is, $P_{\text{data}}(X; W_{\mathcal{G}})$ can be considered the data prior. The data prior $P_{\text{data}}(X; W_{\mathcal{G}})$ is close to the true prior $P_{\mathcal{G}}(X)$ in the sense that $\text{KL}(P_{\text{data}}(X; W_{\mathcal{G}})|P_{\mathcal{G}}(X)) \leq \text{KL}(P_{\text{data}}(Y)|P_{\mathcal{G}}(Y; W_{\mathcal{G}}))$, with the right hand side minimized at the maximum likelihood estimate $\hat{W}_{\mathcal{G}}$, hence the data prior $P_{\text{data}}(X; \hat{W}_{\mathcal{G}})$ at $\hat{W}_{\mathcal{G}}$ should be especially close to the true prior $P_{\mathcal{G}}(X)$. In other words, at $\hat{W}_{\mathcal{G}}$, the posteriors $P_{\mathcal{G}}(X|Y, \hat{W}_{\mathcal{G}})$ of all the data points $Y \sim P_{\text{data}}$ tend to pave the whole prior $P_{\mathcal{G}}(X)$. Based on (9) and (11), in the M-step of EM, for each data point $Y$, $W_{\mathcal{G}}$ seeks to reconstruct $Y$ by $g(X; W_{\mathcal{G}})$ from the inferred latent factors $X \sim P_{\mathcal{G}}(X|Y, W_{\mathcal{G}})$. In other words, the M-step seeks to map all those $X \sim P_{\mathcal{G}}(X|Y, W_{\mathcal{G}})$ to $Y$. Since the posteriors $P_{\mathcal{G}}(X|Y, W_{\mathcal{G}})$ of all $Y$ tend to pave the prior $P_{\mathcal{G}}(X)$, overall, the M-step seeks to map the prior $P_{\mathcal{G}}(X)$ to the whole training data set $P_{\text{data}}$. Of course the mapping from all those $X \sim P_{\mathcal{G}}(X|Y, W_{\mathcal{G}})$ to the observed $Y$ cannot be exact. In fact, $g(X; W_{\mathcal{G}})$ actually maps $P_{\mathcal{G}}(X|Y, W_{\mathcal{G}})$ to a $d$-dimensional patch around the $D$-dimensional $Y$. The local patches for all the observed examples patch up the $d$-dimensional manifold form by the $D$-dimensional observed examples and their interpolations. The EM algorithm is a process of density shifting, so that the mapping of $P_{\text{data}}(X; W_{\mathcal{G}})$ by $g(X; W_{\mathcal{G}})$ shifts towards $P_{\text{data}}$.

In the CoopNets algorithm, we avoid dealing with the posterior distributions for the observed examples $P_{\mathcal{G}}(X|Y; W_{\mathcal{G}})$ and the data prior $P_{\text{data}}(X; W_{\mathcal{G}})$. Instead, we generate directly from the prior distribution $P_{\mathcal{G}}(X)$, instead of using $P_{\mathcal{G}}(X|Y; W_{\mathcal{G}})$ of the observed $Y$ to pave the prior. The density shifting is accomplished by tracking the Langevin revisions made by the descriptor.

## 10. Appendix 4: ConvNet and Alternating back-propagation

Both the descriptor net and the generator net are parametrized by ConvNets. In general, for an $L$-layer net-

work $Y = f(X; W)$, it can be written recursively as

$$X^{(l)} = f_l(W_l X^{(l-1)} + b_l), \tag{25}$$

where $l = 1, ..., L$, $X^{(0)} = X$, $X^{(L)} = f(X; W)$, and $W = (W_l, b_l, l = 1, ..., L)$. $f_l$ is element-wise non-linearity.

The alternating back-propagation algorithms in both the descriptor and the generator involve computations of $\partial f(X; W)/\partial W$ for updating $W$, and $\partial f(X; W)/\partial X$ for Langevin sampling of $X$. Both derivatives can be computed by the chain-rule back-propagation, and they share the computation of

$$\partial X^{(l)}/\partial X^{(l-1)} = f_l^{'}(W_l X^{(l-1)} + b_l) W_l \tag{26}$$

in the chain rule, where $f_l^{'}$ is the derivative matrix of $f_l$. Because $f_l$ is element-wise, $f_l^{'}$ is a diagonal matrix, where the diagonal elements are element-wise derivatives. The computations $\partial X^{(l)}/\partial X^{(l-1)}$ for $l = 1, ..., L$ give us $\partial f/\partial X = \partial X^{(L)}/\partial X^{(0)}$. In addition, $\partial f/\partial W_l = \partial X^{(L)}/\partial X^{(l)} \times \partial X^{(l)}/\partial W_l$. Therefore the computations of $\partial f/\partial X$ are a subset of the computations of $\partial f/\partial W$. In other words, the back-propagation for Langevin dynamics is a by-product of the back-propagation for parameter updating in terms of coding. The two alternating back-propagation algorithms for the two nets are very similar to each other in coding.

For ReLU non-linearity, $f_l^{'} = \mathrm{diag}(1(W_l X^{(l-1)} + b_l > 0)) = \delta_l$, where $\mathrm{diag}(v)$ makes the vector $v$ into a diagonal matrix, and for a vector $v$, $1(v > 0)$ denotes a binary vector whose elements indicate whether the corresponding elements of $v$ are greater than 0 or not. $\delta_l$ is the activation pattern of the ReLU units at layer $l$.

## References


Amari, Shun-ichi and Nagaoka, Hiroshi. *Methods of information geometry*, volume 191. American Mathematical Soc., 2007.

Blum, Avrim and Mitchell, Tom. Combining labeled and unlabeled data with co-training. In *Proceedings of the eleventh annual conference on Computational learning theory*, pp. 92–100. ACM, 1998.

Cover, Thomas M and Thomas, Joy A. *Elements of information theory*. John Wiley & Sons, 2012.

Dai, Jifeng, Lu, Yang, and Wu, Ying Nian. Generative modeling of convolutional neural networks. In *ICLR*, 2015.

Dempster, Arthur P, Laird, Nan M, and Rubin, Donald B. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society. Series B (methodological)*, pp. 1–38, 1977.

Denton, Emily L, Chintala, Soumith, Fergus, Rob, et al. Deep generative image models using a laplacian pyramid of adversarial networks. In *Advances in Neural Information Processing Systems*, pp. 1486–1494, 2015.

Dinh, Laurent, Sohl-Dickstein, Jascha, and Bengio, Samy. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.

Dosovitskiy, E, Springenberg, J. T., and Brox, T. Learning to generate chairs with convolutional neural networks. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.

Girolami, Mark and Calderhead, Ben. Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.

Goodfellow, Ian, Pouget-Abadie, Jean, Mirza, Mehdi, Xu, Bing, Warde-Farley, David, Ozair, Sherjil, Courville, Aaron, and Bengio, Yoshua. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2672–2680, 2014.

Grenander, Ulf and Miller, Michael I. *Pattern theory: from representation to inference*. Oxford University Press, 2007.

Guo, Cheng-En, Zhu, Song-Chun, and Wu, Ying Nian. Modeling visual patterns by integrating descriptive and generative methods. *International Journal of Computer Vision*, 53(1):5–29, 2003.

Han, Tian, Lu, Yang, Zhu, Song-Chun, and Wu, Ying Nian. Alternating back-propagation for generator network. *arXiv preprint arXiv:1606.08571*, 2016.

Hinton, Geoffrey, Vinyals, Oriol, and Dean, Jeff. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.

Hinton, Geoffrey E. Training products of experts by minimizing contrastive divergence. *Neural Computation*, 14 (8):1771–1800, 2002.

Hinton, Geoffrey E, Dayan, Peter, Frey, Brendan J, and Neal, Radford M. The "wake-sleep" algorithm for unsupervised neural networks. *Science*, 268(5214):1158–1161, 1995.

Hinton, Geoffrey E., Osindero, Simon, and Teh, Yee-Whye. A fast learning algorithm for deep belief nets. *Neural Computation*, 18:1527–1554, 2006.

Hopfield, John J. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8): 2554–2558, 1982.

Hyvärinen, Aapo, Karhunen, Juha, and Oja, Erkki. *Independent component analysis*, volume 46. John Wiley & Sons, 2004.

Ioffe, Sergey and Szegedy, Christian. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.

Kim, Taesup and Bengio, Yoshua. Deep directed generative models with energy-based probability estimation. *arXiv preprint arXiv:1606.03439*, 2016.

Kingma, Diederik P. and Welling, Max. Auto-encoding variational bayes. *ICLR*, 2014.

Kirkpatrick, Scott, Gelatt, C Daniel, Vecchi, Mario P, et al. Optimization by simmulated annealing. *science*, 220 (4598):671–680, 1983.

Krizhevsky, Alex, Sutskever, Ilya, and Hinton, Geoffrey E. Imagenet classification with deep convolutional neural networks. In *NIPS*, pp. 1097–1105, 2012.

LeCun, Yann, Bottou, Léon, Bengio, Yoshua, and Haffner, Patrick. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

LeCun, Yann, Chopra, Sumit, Hadsell, Rata, Ranzato, Mare'Aurelio, and Huang, Fu Jie. A tutorial on energy-based learning. In *Predicting Structured Data*. MIT Press, 2006.

Lee, Honglak, Grosse, Roger, Ranganath, Rajesh, and Ng, Andrew Y. Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In *ICML*, pp. 609–616. ACM, 2009.

Liu, Ziwei, Luo, Ping, Wang, Xiaogang, and Tang, Xiaoou. Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3730–3738, 2015.

Lu, Yang, Zhu, Song-Chun, and Wu, Ying Nian. Learning FRAME models using CNN filters. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.

Marinari, Enzo and Parisi, Giorgio. Simulated tempering: a new monte carlo scheme. *EPL (Europhysics Letters)*, 19(6):451, 1992.

Mnih, Andriy and Gregor, Karol. Neural variational inference and learning in belief networks. In *ICML*, 2014.

Neal, Radford M. Mcmc using hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2, 2011.

Ngiam, Jiquan, Chen, Zhenghao, Koh, Pang Wei, and Ng, Andrew Y. Learning deep energy models. In *International Conference on Machine Learning*, 2011.

Olshausen, Bruno A and Field, David J. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision Research*, 37(23):3311–3325, 1997.

Oord, Aaron van den, Kalchbrenner, Nal, and Kavukcuoglu, Koray. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*, 2016.

Pascanu, Razvan, Montufar, Guido, and Bengio, Yoshua. On the number of response regions of deep feed forward networks with piece-wise linear activations. *arXiv preprint arXiv:1312.6098*, 2013.

Radford, Alec, Metz, Luke, and Chintala, Soumith. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.

Rezende, Danilo J., Mohamed, Shakir, and Wierstra, Daan. Stochastic backpropagation and approximate inference in deep generative models. In Jebara, Tony and Xing, Eric P. (eds.), *ICML*, pp. 1278–1286. JMLR Workshop and Conference Proceedings, 2014.

Robbins, Herbert and Monro, Sutton. A stochastic approximation method. *The annals of mathematical statistics*, pp. 400–407, 1951.

Roth, Stefan and Black, Michael J. Fields of experts: A framework for learning image priors. In *CVPR*, volume 2, pp. 860–867. IEEE, 2005.

Rubin, Donald B and Thayer, Dorothy T. Em algorithms for ml factor analysis. *Psychometrika*, 47(1):69–76, 1982.

Salakhutdinov, Ruslan and Hinton, Geoffrey E. Deep boltzmann machines. In *AISTATS*, 2009.

Teh, Yee Whye, Welling, Max, Osindero, Simon, and Hinton, Geoffrey E. Energy-based models for sparse overcomplete representations. *Journal of Machine Learning Research*, 4(Dec):1235–1260, 2003.

Thibodeau-Laufer, Eric, Alain, Guillaume, and Yosinski, Jason. Deep generative stochastic networks trainable by backprop. 2014.

Tu, Zhuowen. Learning generative models via discriminative approaches. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–8. IEEE, 2007.

Vedaldi, A. and Lenc, K. Matconvnet – convolutional neural networks for matlab. In *Proceeding of the ACM Int. Conf. on Multimedia*, 2015.

Welling, Max, Zemel, Richard S, and Hinton, Geoffrey E. Self supervised boosting. In *Advances in neural information processing systems*, pp. 665–672, 2002.

Wu, Ying Nian, Zhu, Song-Chun, and Guo, Cheng-En. From information scaling of natural images to regimes of statistical models. *Quarterly of Applied Mathematics*, 66:81–122, 2008.

Xie, Jianwen, Hu, Wenze, Zhu, Song-Chun, and Wu, Ying Nian. Learning sparse frame models for natural image patterns. *International Journal of Computer Vision*, pp. 1–22, 2014.

Xie, Jianwen, Lu, Yang, Zhu, Song-Chun, and Wu, Ying Nian. A theory of generative convnet. In *ICML*, 2016.

Younes, Laurent. On the convergence of markovian stochastic algorithms with rapidly decreasing ergodicity rates. *Stochastics: An International Journal of Probability and Stochastic Processes*, 65(3-4):177–228, 1999.

Zeiler, Matthew D and Fergus, Rob. Visualizing and understanding convolutional neural networks. In *ECCV*, 2014.

Zhu, Song-Chun. Statistical modeling and conceptualization of visual patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(6):691–712, 2003.

Zhu, Song-Chun, Wu, Ying Nian, and Mumford, David. Minimax entropy principle and its application to texture modeling. *Neural Computation*, 9(8):1627–1660, 1997.